



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 1 130 850 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
05.09.2001 Bulletin 2001/36

(51) Int Cl.7: H04L 12/26

(21) Application number: 01301751.2

(22) Date of filing: 26.02.2001

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

(72) Inventor: Cen, Shanwei
Portland, Oregon 97229 (US)

(74) Representative: Molyneaux, Martyn William
Langner Parry
52-54 High Holborn
London WC1V 6RR (GB)


(30) Priority: 01.03.2000 US 516649

(71) Applicant: Tektronix, Inc.
Beaverton, OR 97077-0001 (US)

(54) Non-intrusive measurement of end-to-end network properties

(57) A method and apparatus for non-intrusive measurement of end-to-end properties of a network flow uses a passive approach. Data units at the input of a network path are sampled by an ingress monitor, which identifies each data unit with a unique signature and generates measurement data. The same data units at the output of the network path are sampled by an egress monitor, which identifies each data unit with a unique

signature and generates measurement data. A data correlator pulls lists of entries from the ingress and egress monitors, each entry having the unique signature and the measurement data for one of the data units, correlates the lists to find entries from the lists that have the same unique signature, and determines from the measurement data for the data units from the two lists that have the same unique signature the end-to-end properties of the network flow.



Description

BACKGROUND OF THE INVENTION

[0001] The present invention relates to testing of networks, and more particularly to a method and apparatus for non-intrusive measurement of end-to-end properties of data transmission through a network path.

[0002] A network provides uni-directional virtual or physical path or channel for data transmission. A *network flow* is a session of data transmission through a network path. In a network flow data units flow from an ingress to an egress of the network path. Example networks are Internet Protocol (IP) networks, Asynchronous Transmission Mode (ATM) networks and circuit-switched, such as telephony, networks. Example data units are IP packets, ATM cells and data bytes (octets). For simplicity the following is described in the environment of IP networks, so the terms *IP networks*, *IP virtual path*, *IP packet* and *packet flow* are used.

[0003] The end-to-end properties of a packet flow may be measured by several metrics, such as throughput, *end-to-end latency*, packet loss, etc. There are two approaches to measuring the metrics. One is active and the other is passive.

[0004] In an active approach packets for monitoring purposes, with special identities and time-stamps, are injected into the ingress of a packet flow. At the egress these special monitoring packets are identified, removed from the flow and time-stamped. Metrics, such as end-to-end latency and packet loss may be measured by analyzing the received monitoring packets. An advantage of the active approach is that it is straightforward and simple, and may be used either online or offline. The main disadvantage of the active approach when used online is that, no matter how small the injected monitoring packet flow is, these monitoring packets always interfere with the operational packet flow. Also, since the network condition is dynamic, off-line measurement might not reflect the actual properties of the packet flow.

[0005] In a passive approach no monitoring packets are injected into the packet flow being monitored. Instead at both the ingress and egress of the packet flow the operational data packets are intercepted without interference to the packet flow itself, and measurement data is generated for each data packet. The measurement data from the two points are correlated, and end-to-end property metrics are derived. In contrast to the active approach the passive approach is non-intrusive. It does not interfere with the operational packet flow in any manner. The cost of this benefit is that the passive approach is much more complex and harder to implement. Due to the reliance on the operational packet flow, the passive approach does not assess the end-to-end properties of a network path before its actual operation.

[0006] A passive approach for ATM network monitoring systems is the WAND project by the University of

Waitato in New Zealand which is disclosed at <http://atm.cs.waikato.ac.nz/wand>. WAND focuses on delay measurement for ATM networks, specifically uses a CRC for data correlation (matching of ATM cells), and matches one cell from the ingress to a list of cells from the egress.

[0007] What is desired is a non-intrusive method of monitoring end-to-end properties of network flows in an efficient way.

BRIEF SUMMARY OF THE INVENTION

[0008] Accordingly the present invention provides a method and apparatus for non-intrusive measurement of end-to-end properties of network flows using a passive approach. An ingress monitor non-intrusively intercepts data units as they enter a network path through a network. Likewise an egress monitor non-intrusively intercepts the same data units as they leave the same network path. Each monitor generates a time stamp for each intercepted data unit using a common clock and derives a unique signature for each data unit such that the same data unit has the same signature at the entry as at the exit. Additionally each monitor counts the number of packets received from the network flow at the ingress and egress respectively. The signature, time stamp and packet counter value form an entry which is retained in an entry queue in each monitor. A data correlator coupled to an out-of-band network to which the monitors also are coupled periodically pulls a list of entries from each monitor and correlates the lists by matching signatures. The time stamps for the entries from the respective lists having the same signature are processed to obtain end-to-end latency of the network path, and the difference in counts over a given time period provides end-to-end packet loss measurements.

[0009] The objects, advantages and other novel features of the present invention are described in the following detailed description when read in conjunction with the appended claims and attached drawing.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWING

[0010] Fig. 1 is an architecture view of a system for non-intrusive monitoring of end-to-end properties of a network path according to the present invention.

[0011] Fig. 2 is a flow chart diagram of a method of non-intrusive monitoring of end-to-end properties of a network path according to the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[0012] Referring now to Fig. 1 an overall architecture of a non-intrusive approach for monitoring end-to-end properties of a network path is shown. To measure the end-to-end properties a monitor 10, 20 is placed at each of the ingress 12 and egress 22 of a packet flow 14 through a network 16. These two monitors 10, 20 are

coupled to a data correlator 18 through an out-of-band network 24 that is independent of and does not interfere with the packet flow 14 being monitored. The two monitors 10, 20 and the data correlator 18 have a common clock 26. The two monitors 10, 20 at the two end points 12, 22 intercept packets from the network flow being monitored, count the data packets received, and extract measurement data from the packets, time stamping the measurements. The data correlator 18 pulls the measurement data from the two monitors 10, 20, correlates them to identify the data from the two monitors belonging to the same packets, and derives measurement results.

[0013] The monitor 10 at the ingress 12 intercepts in real time all the data packets flowing into the path 14 in a non-intrusive way. When a data packet from the data packet flow 14 is intercepted, the monitor 10 generates a time-stamp and extracts a signature of the data packet. The monitor 10 also keeps a count of all packets received so far. The packet signature is a piece of information to be used to differentiate the data packet in question from other data packets in the same packet flow 14. In most cases a special form of Cyclic Redundancy Check (CRC) of the data packet, such as CRC of all fields in a data packet except those that may change hop by hop in an IP packet header, may be used as the packet signature. Other information, such as some fields of the data packet header, may also be extracted as needed. As a result for each data packet the monitor 10 generates an entry composed of <signature, time-stamp, count>, i.e., the packet signature and time-stamp together with a count of the packets received so far. The monitor 10 maintains a queue to hold all recent entries in the incremental order of the time-stamps for retrieval and processing by the data correlator 18. Exactly how long a period of times entries are held depends on several factors, such as the end-to-end packet flow latency, bandwidth and latency of the out-of-band network 24 between the data correlator 18 and monitors 10, 20.

[0014] The monitor 20 at the egress 22 does exactly the same things as the monitor 10 at the ingress 12 on the data packets flowing out of the egress.

[0015] The data correlator 18 samples end-to-end property measurements periodically. During such sampling the data correlator 18 pulls a list of entries from each of the two monitors 10, 20, correlates the two lists and derives end-to-end packet flow property measurements. The data correlator 18 correlates the two lists by matching the packet signatures (S-i; S-e) from the two lists until a subset in one list matches a subset in the other list. When the correlation is done, if one entry in one list matches one entry in the other list, these two entries contain the same packet signature and are considered to be the measurements for the same data packet (P-1). The entry from the ingress monitor 10 (P-i-1: <S-i-1, T-i-1, C-i-1>) contains the time-stamp (T-i-1) when the data packet enters the packet flow, while the

entry from the egress monitor 20 (P-e-1: <S-e-1, T-e-1, C-e-1>) contains the time-stamp (T-e-1) when the data packet leaves the packet flow, where S-i-1 = S-e-1. The difference between these two time-stamps is the end-to-end latency for the data packet to flow through the network path 14. The frequency with which the data correlator 18 pulls data from the monitors 10, 20 and the amount of data it pulls each time depends on the specific configuration in use, the amount of available bandwidth of the out-of-band network 24, and other relevant factors.

[0016] At a later time n the data correlator 18 identifies another pair of matched entries (P-i-n: S-i-n, T-i-n, C-i-n), (P-e-n: S-e-n, T-e-n, C-e-n) where S-i-n = S-e-n, then $((C-i-n - C-i-1) - (C-e-n - C-e-1))$ is the number of packets lost between P-1 and P-n. The packet loss ratio is $((C-i-n - C-i-1) - (C-e-n - C-e-1)) / ((C-i-n - C-i-1) - (C-e-n - C-e-1))$ and the packet loss rate is $((C-i-n - C-i-1) - (C-e-n - C-e-1)) / (T-i-n - T-i-1)$.

[0017] The following further illustrates the measurement data correlation. Suppose at time T the data correlator 18 pulls from the monitors 10, 20 all entries between T_s and T, where $T_s < T$. The list from the ingress monitor 10 is: List-i: (P-i-1: <S-i-1, T-i-1, C-i-1>, ..., P-i-m: <S-i-m, T-i-m, C-i-m>) where (P-i-1: <S-i-1, T-i-1, C-i-1>) is an entry with packet signature S-i-1, time-stamp T-i-1 and count C-i-1. The list from the egress monitor is: List-e: (P-e-1: <S-e-1, T-e-1, C-e-1>, ..., P-e-n: <S-e-n, T-e-n, C-e-n>) Suppose that the sub-list (P-i-x1, ..., P-i-xk) of List-i matches the sub-list (P-e-y1, ..., P-e-yk) of List-e, i.e., S-i-xj = S-e-yj for j from 1 to k. In that case S-i-xj and S-e-yj are considered to be the packet signatures of the same data packet P-j, and T-i-xj and T-e-yj are the times when the data packet P-j passed the ingress 12 and egress 22 respectively. Then T-e-yj - T-i-xj is the end-to-end latency for the data packet to flow through the network path 14. By calculating the measurements for all the data packets in the matched sub-lists, a set of latency measurements is derived.

[0019] Similarly the packet loss between any two data packets may be calculated from their packet count values, as illustrated above.

[0020] The latency measurements on individual data packets generated in the above process may be used for analysis of various short-term and long-term latencies of network flows. The latency measurement samples taken over a long period of time (hours, days, weeks or months) may be used to generate statistics, such as percentile, histogram and distribution, for analysis of the long-term distributions and trends in the latency of network flows. Each list of latency measurements that the data correlator 18 gets covers a short period of time and may be used to derive short-term latency jitter. By sampling lists of latency measurements periodically the long-term trends of latency jitter may be analyzed.

[0021] Similarly based on the individual packet loss measurements short-term and long-term packet loss

properties may be analyzed.

[0022] In the above discussion for purposes of illustration it is assumed that the data correlator 18 periodically pulls measurement packet lists from the ingress and egress monitors 10, 20. In an actual application the ingress and egress monitors 10, 20 may push the measurement packets to the data correlator 18. In fact the push model may be more efficient in some cases.

[0023] The architecture of Fig. 1 is conceptual. In actual implementations in order to reduce the bandwidth requirements of the out-of band network 24 the data correlator 18 function may be located in one of the two monitors 10, 20. The combined data processor/monitor implementation has at least two advantages: (a) there is a reduced network bandwidth requirement; and (b) it is feasible for the data correlator 18 to get a much longer list of data from the resident monitor 10, 20, thus increasing the search space during data correlation. On the other hand a physically separate data correlator 18 may easily be shared by multiple pairs of monitors 10, 20 monitoring different packet flows. This data correlator 18 may also serve as a centralized location for measurement network configuration and management. In the case where a large number of packet flows are to be monitored simultaneously, more than one data correlator 18 may be needed to distribute the workload among them.

[0024] Subsampling of measurement packets may be necessary if the monitors 10, 20, the out-of-band network 24 or the data processor 18 cannot capture, transport or process the measurement data of all the data packets in the packet flow. With measurement packet subsampling the monitors 10, 20 still count all packets, but only generate measurement data of a selective subset of the data packets in the packet flow, and the end-to-end properties are derived from the measurement data of this subset of data packets. There are two requirements a subsampling scheme has to satisfy. One requirement is that the ingress and egress monitors 10, 20 select the same subset of data packets. Otherwise if the two measurement data lists from the two end-points 12, 22 are for different sets of data packets, it becomes impossible to derive the end-to-end properties. The other requirement is that the sampling be periodical or random additive based on some statistical distribution. Signature bit masking is a subsampling scheme that satisfies both requirements. With signature bit masking a data packet is selected only if the least significant part of its signature has a specific given value. For example to subsample one-eighth of the data packets, select only those data packets whose signature has 000 as its least significant three bits. If a data packet is selected by the ingress monitor 10, it also is selected by the egress monitor 20 since the packet signature does not change inside the packet flow 14. When CRC is used for the signature, since CRC is well randomized, signature bit masking becomes a very good form of geometric sampling in which a packet is selected with a fixed probability,

which is a form of random additive sampling.

[0025] The uniqueness of the packet signature affects the reliability of the data correlation by the data correlator 18. Different data packets may have the same content and thus the same CRC. Even data packets with different contents may have the same CRC. Non-uniqueness of packet signatures introduces ambiguity in the list correlation. Carefully selecting the fields to be included in the special CRC for the packet signature helps to improve uniqueness. For example since the packet identification field in an IP header contains the value of a counter which normally increments by one upon sending of each IP packet by a host, inclusion of this field in the packet signature CRC may make the packet signatures of multiple packets unique even if they contain the same payload. There may be cases when there are more than one way the measurement data lists from the two monitors 10, 20 match against each other. If this ambiguity occurs, there is no way to tell reliably which match is correct. As a result it is not possible to derive latency measurements from the two lists in question. Increasing the number of entries in the list of measurement data pulled from the monitors 10, 20 helps to reduce the probability of ambiguity, but requires more bandwidth from the out-of-band network 24.

[0026] Some of the properties of IP networks made data correlation more complicated. IP packets may not only be dropped, they may also be duplicated or fragmented by the network 16. They also may be delivered out of order. Depending upon the configuration of the network 16, some of these problems may not exist.

[0027] When the present invention is applied to ATM networks, some form of CRC may still be used as the cell signature. When applied to circuit-switched networks in which the data units are bytes, the whole byte itself may be used in place of its signature.

[0028] Thus the present invention provides non-intrusive measurements of end-to-end properties of network flows by sampling the data packets at both the ingress and egress of the network flow being measured, assigning a unique signature and time-stamp using a common system clock for each packet as well as a count of packets received so far, correlating lists of packets from both ingress and egress based on the unique signature, and using the time-stamps and counts for the same packet at ingress and egress obtaining measurements of the desired end-to-end properties of the network flow.

50 Claims


1. A method of non-intrusive measurement of end-to-end properties of a network flow, the network flow being a session of data units flowing from an ingress to an egress of a network path, comprising the steps of:

sampling data units at the ingress using an in-

- gress monitor;
 generating a unique signature and input measurement data for each sampled data unit at the ingress monitor;
 sampling the data units from the network flow at the egress using an egress monitor;
 generating the unique signature and output measurement data for each sampled data unit at the egress monitor;
 correlating entries from each of the ingress and egress monitors, the entries being in the form of a list for each monitor and each entry having the unique signature and associated measurement data for one of the data units, by matching the unique signatures from entries between the two lists; and
 determining from the input and output measurement data for the matched entries the end-to-end properties of the network flow.
2. The method as recited in claim 1 wherein the input and output measurement data are timestamps indicating the time of receipt of the data units at the ingress and egress monitors, the timestamps being derived from a common clock.
 3. The method as recited in claims 1 or 2 wherein the input and output measurement data are counts of the numbers of data units received by the ingress and egress monitors so far.
 4. The method as recited in claim 1 further comprising the step of subsampling the data units prior to the generating steps such that the ingress and egress monitors select the same subsampled data units with a random additive distribution.
 5. The method as recited in claim 4 wherein the subsampling step comprises the step of masking of the unique signature such that only if the least significant portion of the unique signature has a specific given value is the data unit selected by the ingress and egress monitors.
 6. A system for non-intrusive measurement of end-to-end properties of a network flow comprising:
 - an ingress monitor coupled to the input of the network flow for intercepting data units flowing through a network path, the ingress monitor identifying a unique signature and generating input measurement data for each data unit;
 - an egress monitor coupled to the output of the network flow for intercepting the data units flowing through the network path, the egress monitor identifying the unique signature and generating output measurement data for each data unit; and

a data correlator coupled to an out-of-band network to which the ingress and egress monitors are also coupled, the data processor pulling lists of entries from the ingress and egress monitors, the entries containing the unique signature and measurement data for each data unit, correlating the lists of entries based on the unique signatures, and determining from the input and output measurement data of data units having the same unique signatures the end-to-end properties for the network flow.

7. The system as recited in claim 6 wherein the input and output measurement data are timestamps for the data units derived from a common clock indicating the time of arrival at the ingress and egress monitors respectively.
8. The system as recited in claims 5 or 6 wherein the input and output measurement data are counts of the number of data units received at the ingress and egress monitors respectively so far.
9. The system as recited in claim 5 wherein the ingress and egress monitors subsample the data units prior to generating measurement data, the data units selected being the same at the ingress and egress monitors with a random additive distribution.
10. The system as recited in claim 9 wherein the data units are subsampled based upon a given value for a least significant portion of the unique signature.






FIG.2

THIS PAGE BLANK (USPTO)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) **EP 1 130 850 A3**

(12) **EUROPEAN PATENT APPLICATION**

(88) Date of publication A3:
25.06.2003 Bulletin 2003/26

(51) Int Cl.7: **H04L 12/26, H04L 12/24**

(43) Date of publication A2:
05.09.2001 Bulletin 2001/36

(21) Application number: **01301751.2**

(22) Date of filing: **26.02.2001**

(84) Designated Contracting States:
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR**
Designated Extension States:
AL LT LV MK RO SI

(72) Inventor: **Cen, Shanwei**
Portland, Oregon 97229 (US)

(74) Representative: **Molyneaux, Martyn William**
Wildman, Harrold, Allen & Dixon
11th Floor, Tower 3,
Clements Inn,
London WC2A 2AZ (GB)


(30) Priority: **01.03.2000 US 516649**

(71) Applicant: **Tektronix, Inc.**
Beaverton, OR 97077-0001 (US)

(54) **Non-intrusive measurement of end-to-end network properties**

(57) A method and apparatus for non-intrusive measurement of end-to-end properties of a network flow uses a passive approach. Data units at the input of a network path are sampled by an ingress monitor, which identifies each data unit with a unique signature and generates measurement data. The same data units at the output of the network path are sampled by an egress monitor, which identifies each data unit with a unique

signature and generates measurement data. A data correlator pulls lists of entries from the ingress and egress monitors, each entry having the unique signature and the measurement data for one of the data units, correlates the lists to find entries from the lists that have the same unique signature, and determines from the measurement data for the data units from the two lists that have the same unique signature the end-to-end properties of the network flow.



EP 1 130 850 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 30 1751

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	GRAHAM I ET AL: "COMPARATIVE MEASUREMENT OF QOS ON THE TRANS-PACIFIC INTERNET" PROCEEDINGS OF THE SPIE, SPIE, BELLINGHAM, VA, US, vol. 3529, 2 November 1998 (1998-11-02), pages 289-294, XP008008877	1-3,6-8	H04L12/26 H04L12/24
A	* the whole document *	4,5,9,10	
X	EP 0 948 165 A (HEWLETT PACKARD CO) 6 October 1999 (1999-10-06)	1,6	
A	* abstract *	2-5,7-10	
	* claims 1,2,5,8,13 *		
	* paragraph [0004] *		
	* paragraphs [0013]-[0016] *		
	* paragraph [0030] *		
	* column 16, line 11-18 *		
A	EP 0 716 525 A (AT & T CORP) 12 June 1996 (1996-06-12)	1-10	
	* abstract *		
	* claims 1-3,10 *		
	* figures 1,3 *		
	* page 2, line 24 - page 2, line 59 *		
	* page 3, line 53 - page 4, line 22 *		
A	US 5 764 912 A (ROSBOROUGH JAMES M) 9 June 1998 (1998-06-09)	1-10	
	* abstract *		
	* claims 1,2,15,16 *		
	* column 1, line 45 - column 2, line 21 *		
	* column 3, line 7-9 *		
	* column 3, line 61 - column 4, line 10 *		
	* column 4, line 47-62 *		
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 7 May 2003	Examiner Cichra, M
CATEGORY OF CITED DOCUMENTS		T: theory or principle underlying the invention E: earlier patent document, but published on, or after the filing date O: document cited in the application L: document cited for other reasons &: member of the same patent family, corresponding document	
X: particularly relevant if taken alone Y: particularly relevant if combined with another document of the same category A: technological background O: non-written disclosure P: intermediate document			

EPO FORM 1503 03.02 (P04C01)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 01 30 1751

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

07-05-2003

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0948165 A	06-10-1999	EP 0948163 A1	06-10-1999
		EP 0948164 A1	06-10-1999
		EP 0948165 A1	06-10-1999
EP 0716525 A	12-06-1996	US 5570346 A	29-10-1996
		CA 2160616 A1	09-06-1996
		EP 0716525 A2	12-06-1996
		JP 8340353 A	24-12-1996
US 5764912 A	09-06-1998	US 5781449 A	14-07-1998
		AU 699693 B2	10-12-1998
		AU 6505396 A	05-03-1997
		BR 9610130 A	20-06-2000
		CA 2229022 A1	20-02-1997
		CN 1197521 A	28-10-1998
		EP 0871938 A1	21-10-1998
		JP 11510660 T	14-09-1999
		WO 9706494 A1	20-02-1997
		US 2002110091 A1	15-08-2002
		US 6493754 B1	10-12-2002
		US 6308211 B1	23-10-2001
		US 5838920 A	17-11-1998

EPO FORM P0458

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

THIS PAGE BLANK (USPTO)

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☐ ~~FADED TEXT OR DRAWING~~
- ☒ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

THIS PAGE BLANK (USPTO)